

## Good News, Bad News: A Sentiment Analysis of the 2016 Election Russian Facebook Ads

GERMAN ALVAREZ

JAEWON CHOI

SHARON STROVER<sup>1</sup>

University of Texas at Austin, USA

In the wake of the 2016 U.S. presidential election, the notion of “fake news” and Russian election interference became somewhat interchangeable. We argue that “fake news” insufficiently describes the Russian disinformation campaign. Our analysis of Facebook ad texts shows that they incorporate emotional appeals differently at unique moments in the political campaign of 2016. We use sentiment analysis to demonstrate the use of positive and negative emotional messages. Sentiment ratings dipped to more negative scores before the November 2016 election and rose to new positive heights after the election. This provides some evidence that the varying uses of positive sentiment and negative sentiment may have been strategic.

*Keywords: Facebook, disinformation, Russian propaganda, emotional appeal*

On December 8, 2016, Hillary Clinton made a speech decrying the role of “fake news” on social media. She warned her audience that fake news had real-world consequences, possibly a reference to her recent loss in the 2016 presidential election. The following month, President Trump, in a curt explanation on why he would not call on a CNN reporter at a press conference, said it was because they were “fake news.” It is arguably around this time that *fake news* became a permanent fixture in the public lexicon (Silverman & Alexander, 2016).

The imprecision and ambiguity of the phrase *fake news* enable the term to mean different things for different people. Sometimes it is used interchangeably with “misinformation” or “disinformation.” Some simply call it propaganda. President Trump now uses the term like a slur with which he insults negative news coverage of himself or his administration. What is clear here is that fake news has been repurposed as a rhetorical device for political actors. We know more about Russia’s 2016 interference as a result of the Cambridge Analytica

---

German Alvarez: gaalvarez@utexas.edu

Jaewon Choi: jaewonrchoi@utmail.utexas.edu

Sharon Strover: sharon.strover@austin.utexas.edu

Date submitted: 2019–02–04

<sup>1</sup> We greatly appreciate and acknowledge the support of the Good Systems program at the University of Texas at Austin.

Copyright © 2020 (German Alvarez, Jaewon Choi, and Sharon Strover). Licensed under the Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd). Available at <http://ijoc.org>.

scandal and subsequent Congressional investigations in 2018, as well as from the release of Special Counsel Robert Mueller's (2019) report. Jamieson's (2018) phrase *viral deception* captures both the distribution mode and the deceptive intent of the Russian-produced materials, and she helpfully positions the Facebook ads in the context of years of communication research that has repeatedly addressed the role of textual and visual appeals in campaign messages. She hints at something that Dahlgren (2018) elaborates in his discussion of rationality and emotionality in social media. Both underscore the position that ignoring emotion in democratic participatory modes is counterproductive. Dahlgren goes further in noting that emotional expression is easily encouraged in the online sphere, but it may undercut actual functional participation in politics.

An overview of the Russian interference campaign suggests that fake news and even the bland term *interference* may be insufficient to fully grasp the Russian protocols. We argue that the combined affordances and platforms of social media and their abilities to quickly capture and convey emotion mean that we must examine materials such as the Russian Facebook ads in terms that can grapple with emotional intent or impact. The goal here is to analyze the Russian-inspired ads to investigate how they convey emotional sentiment amid timeworn themes that frame and prime salient and tumultuous political and cultural issues. Our unique contribution is to move beyond matters of facticity in such ads to highlight ways that emotional language on social media platforms is related to engagement practices. Those in turn may interfere with productive political participation.

The Mueller (2019) report concludes that social media platforms such as Facebook and Twitter had insufficient monitoring and content-tracking practices in place to detect content that was produced by groups or individuals misrepresenting themselves and their interests. The vulnerability of these platforms to trolls, bots, and what appears to be foreign countries attempting to interfere with domestic information flows has now become clear, and the platforms themselves have sought assistance from the research community on how they might improve their operations (Kelly, 2019).

The extent of the Russian "meddling" with U.S.-directed information is more apparent from the vantage point of 2020, even as security personnel warn that the same sorts of meddling appear to be in play for the 2020 elections (Goldman, Barnes, Haberman, & Fandos, 2020). A U.S. Congressional inquiry into the 2016 meddling found that 30.4 million Facebook posts linked to Russia were shared by users (Roose, 2018). The original ads were publicly released in May 2018 by Congressional committees investigating Russian influence (Guynn, Weise, & Kelley, 2018), and were followed up by later releases from Twitter (10.4 million tweets from Russia's Internet Research Agency [IRA]), YouTube (1,107 videos posted by the IRA), and Instagram.

A notable quality stands out in these social media-based materials. The ads indeed conveyed some false information, but they also placed emotional appeals in people's news feeds that had little to do with conventional truth or lies. Those ads included invitations to Facebook groups centered on divisive topics for the American public, the most numerous of them focusing on race, police brutality, nationalism, immigration, and gun rights (Kim et al., 2018). Four of those topics also were among the top-10 voting issues for the 2016 campaign (Pew Research Center, 2016), and those same subjects frequently catalyze economic, racial,

or cultural unrest or unease. An overview of the IRA's strategy<sup>2</sup> indicates a sophisticated campaign to sow discord among the U.S. public by focusing on these and related topics. These particular appeals—emotional, opinion-laden, and invoking hot-button topics—are undersold by statements such as “Russian fake news helped Trump win.”

From this perspective we analyze the textual content in 3,519 advertisements purchased by Russia's IRA and released by Facebook to the U.S. Congress. We investigate the nature of the Facebook advertisements' use of emotion to generate attention, investigating how positive and negative sentiment might have been strategic. The results underscore the temporal flows of emotion in the Russian ad texts and challenge the efficacy of evaluating such messages simply in terms of truth or falsity.

### Literature Review

The long history of propaganda research in the communication field has yielded an understanding of how persuasion occurs, and it is highly relevant to any discussion of the content and possible influence of Facebook ads (Eagly & Chaiken, 1993; Katz & Lazarsfeld, 1955; Lasswell, 1927; Popkin, 1994). That research has investigated notions of people's response to certain images, terms, and arguments, the effects accruing to conditions of message exposure and influence of sources. The way a particular category of messages—*disinformation*—operates is pertinent to the current investigation. The research literature reviewed distinguishes the nature of content associated with social media-based disinformation and examines its cognitive and emotional information-processing characteristics. Recent research on social media and Facebook in particular notes their memetic qualities, including their brevity (Facebook ads are a single page), their penchant for illustrating powerful images and text, their ability to capitalize on recognizable memes, their availability for liking and sharing, and of course their ability to network people. Those medium qualities intersect certain types of political address and participation.

### *Disinformation, Cognitive Processing, and Fake News*

Academics and journalists alike argue that they should no longer use the term *fake news* (Gelfert, 2018; Kramer, Guillory, & Hancock, 2014; Van Duyn & Collier, 2018). Although the actual information deployed in the Facebook ads was at times false, that was not its defining characteristic.<sup>3</sup> The term is inaccurate and its use confuses people (Van Duyn & Collier, 2018). *Disinformation* more precisely can be defined as politically motivated and explicitly directed to depress citizen engagement and genuine political participation (Gelfert, 2018; Jackson, 2017). Disinformation resembles propaganda in that the goal is to incite citizens to action and dampen information processing characterized by deliberation and reflection (Stanley, 2015). The goal of propaganda is to subvert rational thinking, and effective propaganda encourages people to move as a collective, motivated by emotions that disregard evidence and substitute

---

<sup>2</sup> The term *strategy* is used to suggest that there is a formula behind the construction of the ad content. Jamieson (2018) unpacks the sequence of what she calls “the Russian courtship” (p. 119) by highlighting the early stage of benign appeals to signal shared interests, followed by extreme statements.

<sup>3</sup> Popular media have highlighted the obvious grammatical and syntactical mistakes in the Facebook ads. Those slips offer a clue to the ads' provenance, but the fact remains that they were often widely circulated nonetheless.

emotional intensity for mindful and thoughtful response. Whether or not the information is true or false is secondary to whether it can sow the desired outcomes among the public. Disinformation shares that goal.

From a communication research perspective, disinformation aims to trigger information-processing mechanisms such as priming, confirmation bias, and affective arousal that likely prompt people to like, share, and increase the reach of messages within an online platform context. In her 2018 book *Cyberwar*, Jamieson ably dissects the potential priming impact of the Facebook ads' invocation of contentious issues such as immigration, police brutality, and attacks on the flag, among other themes. Those types of issues and appeals, especially in their social media incarnation, emphasize emotional content. The mode of spreading disinformation exploits the algorithms on social media that tend to bias sensational and evocative information (Jamieson, 2018; Kramer et al., 2014), enabling a "Pearl Harbor of the social media age" (Roose, 2018, p. 2).

### ***Emotion, Heuristic Processing, and Ads***

The dual processing model of persuasion explains how people process information and form attitudes, contrasting *systematic* or cognitive modes with *heuristic* or only weakly cognitive mechanisms (Chaiken & Ledgerwood, 2012). This literature reflects on the rapid-fire, heavily visual, and slogan-dependent content of visual ads such as those on Facebook as opposed to lengthy written debates requiring time, attention, and intellectual engagement. Beyond demonstrating that there are differences in how certain types of information are processed, this literature underscores the widespread practice of quick and often biased cognitive processing (Elliott, 2008; Kahneman, 2011). In the heuristic process, people use comparatively less cognitive effort and instead rely on more easily accessible information such as source identity and status or cues that can quickly elicit affective associations (Chaiken, 1980; Chaiken & Ledgerwood, 2012; Petty & Cacioppo, 1984).

The meme-like qualities and brevity associated with Facebook posts, particularly with image-based posts, are amenable to such processing. For instance, Buck, Chaudhuri, Georgson, and Kowta (1995) and Chaudhuri and Buck (1995) show that digital and more visually oriented media trigger peripheral cognitive responses, whereas relatively text-heavy traditional print advertisement leads to more central cognitive responses. Kahneman's (2011) basic argument in *Thinking Fast and Slow* underscores the pervasiveness of "fast thinking" and the intuitive and emotional biases that can be exacerbated by the speed of digital content encounters. The design of Facebook ads makes the spread of false or contentious beliefs not simply an unfortunate side effect, but rather the intended function, operating through emotional appeals. Emotions are a core component for how people see the world (Marcus, 2002), and there is considerable support for the notion that mood, feelings, and emotions shape public opinion (Brader, 2006; Jamieson, 2018).

A common mistake, however, in earlier research on emotions and politics was to treat the affective process as a suppressor of rational thinking (Brader, 2006). This line of thinking argues that if people were more rational, then they could subvert the effect of the Russian Facebook ads. This is, however, an antiquated way of understanding the affective and cognitive aspects of human information intake. The two dimensions, affective and cognitive, often work together (Kahneman, 2011). Other researchers argue similarly that motivated reasoning is a dynamic process that involves both affective and cognitive considerations (Kunda, 1990; Nyhan & Reifler, 2010; Taber & Lodge, 2006; Weeks, 2015). Lodge and

Taber's (2000) early theory of affective cognitive processes for motivated reasoning argues that all social concepts and information are tied to an affect, echoing Abelson's (1963) earlier work affirming that once a cognitive node is activated, so too is its related affect. The premise is that when people encounter new information, they revise the affective tags attached to their conceptual memory (Anderson & Hubert, 1963; Clore & Isbell, 2001; Schwarz, 2011). Hence, emotional appeals can invoke affective and cognitive processes in ways that may be logical but not necessarily "rational."

To be effective, the Russian-placed advertisements follow typical campaign advertising formats. That is, political campaign ads classically involve short and simple reasons that emphasize why someone should support a candidate or issue (i.e., the rationale). Coupled with these messages are images, text, and music that evoke an emotional appeal, which Brader (2006) argues impact how the messages in the ad are evaluated (i.e., the affective component). The text in political ads explicitly conveys meaning and may convey emotion, and the imagery and tone enhance how the message is received. For example, a Facebook ad with the message "police brutality is out of control" is enhanced with pictures of helpless African Americans under attack by police. It is in this way that the message and the images work together to deliver the message.

The impact of emotional ad appeals, however, is highly dependent on what emotion is evoked (Carver & Harmon-Jones, 2009; Gervais, 2017; Tannenbaum et al., 2015; Weeks, 2015). For example, research on political ads with a positive message, such as "enthusiastic" or "feel-good" ads, can motivate people toward in-group favoritism (Tannenbaum et al., 2015; Weeks, 2015). As such, these positive ads have shown effectiveness at reinforcing existing loyalties. Negative political ads have been shown to heighten partisanship and out-group antagonism (Gervais, 2017; Weeks, 2015). In fact, anger can motivate partisan, motivated evaluation of misinformation (Weeks, 2015).

Whether campaigns use positive or negative ads is related to several factors including competitiveness of the race (Druckman, Kifer, & Parkin, 2010), ad sponsorship (Fowler, Ridout, & Franz, 2016), and media platform (i.e., television or online; Borah, Fowler, & Ridout, 2018). As it relates to our research, scholars have found that online political ads tend to be more positive compared with traditional television campaign ads (Bekafigo & Pingley, 2015; Borah et al., 2018). Diffusion may gain power with more positive online content. Positive content spreads more quickly and thoroughly online (Berger & Milkman, 2012). The virality of social media content is driven by physiological arousal, and social media content that evokes high positive or negative physiological arousal is more viral. In addition, positive content is related to higher arousal, making it the most effective way to spread content throughout social media. When Brady, Wills, Jost, Tucker, and Van Bavel (2017) examined 563,312 Twitter posts on gun control, same-sex marriage, and climate change, their results showed that people are 20% likelier to share or interact with emotionally evocative social media information. Facebook's own sponsored research on "massive-scale emotional contagion" suggests that emotional states are transferable through Facebook itself (Facebook, n.d.). Using actual news feed content that expressed positive or negative posts, Kramer and colleagues (2014) found that emotions expressed by contacts on Facebook affect users' emotions. (The research procedures in this study were later assessed as highly controversial.) Finally, Cyr, Head, Lim, and Stibe (2018) examined how message argument quality and other heuristic factors such as website design aspects correspond to individuals' issue involvement, which subsequently can lead to attitude change. Their study

demonstrates how both central and heuristic cues function together when shaping attitudes on certain issues, with heuristic responses cued to visual materials.

The potential for social endorsements, such as commenting, liking, or sharing ads of political advertisements on social media, can affect how people evaluate the ads (Borah et al., 2018; Thorson, Vraga, & Ekdale, 2010). Research indicates that "sharing" and "liking" social media content can strengthen social bonds and increase group solidarity and social capital (Pi, Chou, & Liao, 2013; Placencia & Lower, 2013). In addition, posts with positive appeals using humor and enthusiasm are the most likely to be shared (Borah, 2016).

Bakir and McStay (2018) suggest that the Russian Facebook disinformation content used more emotional appeals instead of false or made-up information. The evocative nature of the IRA's Facebook ad content may have incited social conflict around divisive societal issues of the day, namely immigration and racial bias in policing in the lead-up to 2016 U.S. presidential election, a speculation also emerging from the work of Ribeiro, Ufmg, Henrique, Gummadi, and Redmiles (2018). These scholars tested some of the Russian ad content effects with a selected sample. They conjecture that emotional appeals help to get better attention and also inspire group emotional behaviors within social networks, labeling the phenomenon *socially divisive advertising* and echoing results highlighted by Iyengar, Sood, and Lelkes (2012). These researchers found that the use of affective polarization increased throughout the 2008 election campaign, and that it was more prominent in battleground states, a finding later reaffirmed as growing in the century's second decade (e.g., Iyengar, Lelkes, Levendusky, Malhotra, & Westwood, 2019). More important, they found that the volume of negative advertising was positively related to affective polarization. Rooted in social identity theory, affective polarization posits that when social identities are made salient, people tend to put greater weight on their identity-relevant attitudes and favor their in-group (Hogg, Terry, & White, 1995; Tajfel, 1969; Tajfel & Turner, 1979). In-group favoritism is often accompanied by a negative bias toward the out-group. In the case of the Russian advertisements, it is possible that positive sentiments for social groups were made more salient initially and then later threatened or challenged with negative sentiments, the net result being antagonism across various groups.

### ***Empirical Studies on Facebook Political Advertisements***

Given that Facebook has become an accepted news distribution and circulation vehicle, people are used to seeing what loosely could be called "informational material" there. Social media platforms have become major venues for Americans to seek information. One recent Pew Research Center report (Matsa & Shearer, 2018) showed that approximately 68% of American adults indicated that they at least occasionally obtain news on social media; among different available social media platforms, Facebook is most commonly used for news consumption. About 43% of Americans who consume news from social media platforms get it from Facebook (Matsa & Shearer, 2018). In addition to the elevated significance of social media as a main information source, technological factors such as native advertising and microtargeting enable seamless, highly targeted content from these sources (Kim et al., 2018). Such technological affordances make social media platforms such as Facebook attractive places to disseminate tailored political messages during sensitive political events such as presidential elections.

Facebook's ability to target potential recipients by geography and certain demographic characteristics also is highly desirable. Facebook works to balance both the advertiser's goals and the Facebook user's experience (Facebook, n.d.). For advertisers, the system works to reach their target audience. Users want their visit to Facebook to be positive and relevant. To achieve this balance, Facebook uses an automated auction system. However, unlike the "highest bid wins" method, Facebook uses an automated auction system that involves three factors: the bid, estimated action rates, and ad quality and relevance. Advertisers bid well when they are able to make their ad "relevant" to users; that is, Facebook generates improved placements when there is higher engagement. Facebook can estimate an ad's "action rate," the likelihood a user will take action once seeing the ad. The estimate is determined by the ad's historical performance on Facebook and the target audience's past actions. Finally, ad quality and relevance are determined by looking at user feedback toward an ad. Negative feedback for an ad will decrease its total value. If positive emotions prompt more engagement, relevance scores would be improved. This presents another factor in understanding how sentiment can figure into the construction of the Russian-placed ads.

The reach and potential for covert sources to place content affecting users via Facebook have compelled researchers to try to understand where fake news and disinformation come into play in political matters. Analyzing advertisement content and impressions of 783 unique Facebook political advertisements during the 2017 United Kingdom general election campaign, Anstead, Magalhaes, Stupart, and Tambini (2018) found that political advertisements on Facebook resembled the issues and strategies of the wider campaign while allowing certain political parties to more effectively target their constituencies. Whereas the ability to microtarget audiences might be useful to maximize the effectiveness of political campaigns, it could potentially yield harmful implications, creating "filter bubbles" for individuals and cultivating populism (Groshek & Koc-Michalska, 2017). Benkler, Faris, and Roberts (2018) found evidence that Facebook circulates false content and hyperpartisan disinformation more often than Twitter. Furthermore, nefarious tampering can be attempted by malicious domestic as well as foreign institutions (Dwoskin & Romm, 2018). Despite Facebook's stated commitment to cooperate fully in mounting investigations, a severe lack of transparency of these social media platforms' targeting practices has long been a concern for researchers (Kreiss & Howard, 2010).

In part because of Facebook's opacity, only a handful of studies have managed to empirically examine the dynamics of foreign influence on the U.S. political agenda through social media. Some researchers have documented the targets behind the ads, and some have chronicled the topics invoked in the ads. Investigating paid ads on Facebook during the 2016 U.S. presidential election campaign, Kim et al. (2018) identified groups behind divisive issue campaigns on Facebook and the targets for such messages. The researchers found that approximately half of 228 total groups behind the Facebook issue campaigns in the 2017 election fell into a "suspicious group category," referring to sites either taken down by Facebook or showing minimum activity after the election day. About 17% of these suspicious group were found to be Russian-linked groups. An analysis of Russia's IRA's 3,519 Facebook and Instagram advertisements released by the House Permanent Select Committee on Intelligence found that the IRA created 73 different advertisement sponsors and groups on Facebook, which in turn connected to more than 100 groups spreading beyond Facebook-related social media platforms to venues including Google, YouTube, MeetUp, and Twitter (Kim, 2018). Kim et al. (2018) and Ribeiro et al. (2018) reveal that advertisements specifically addressed major campaign issues, with race issue ads the most frequently and prominently distributed.

Furthermore, approximately 30% of total IRA group ads mentioned elections and voting, and about 8% contained mentions of candidate and party, corroborating the idea that IRA groups' intended to influence the election (Kim, 2018), a point also emphasized by Jamieson (2018).

Between 2015 and 2017, the IRA increased its ad purchases from 617 in 2015 to 1,867 in 2016 and 1,033 in 2017. Nearly 25% of the landing pages spent more than \$100, 26.8% of the pages received more than 1,000 clicks, and approximately 36.1% had more than 10,000 impressions.<sup>4</sup> Of the 3,500 ads in the data set released by Congress, more than 50% made explicit references to race (Penzenstadler, Heath, & Guynn, 2018). Twenty-five percent of the ads were related to crime and police with a racial connotation, many involving police brutality or Black Lives Matter and social justice issues. In all, divisive racial ad buys averaged about 44 per month from 2015 through the summer of 2016 before seeing a significant increase in the run-up to the election day. Between September and November 2016, the number of race-related spots rose to 400. An additional 900 were posted after the November election through May 2017. This temporal aspect raises questions as to whether emotional sentiment in the ads changed over time as well.

Finally, the Russian advertisements were carefully crafted to target specific demographic profiles. We know that low-income populations were targeted with immigration- and racial conflict-related advertisements, middle-income populations were shown ads on nationalism, and the majority of immigration and nationalism issue advertisements were targeted toward Whites compared with other racial/ethnic groups (Kim et al., 2018). Incorporating additional survey data on representative users' tendency to report, approve, and identify false claims in the ads, Ribeiro et al. (2018) found that IRA ads appeared to target people who were more likely to believe and approve the advertisement content.

Estimates by the intelligence community as conveyed to Congress indicate that more than 11.4 million American users were exposed to those advertisements as the Facebook ads were further shared with other social groups. The social network strategy was to get users to "like" one of 470 IRA-constructed Facebook pages. According to Facebook, once liked, Facebook users automatically follow the page, which means that they may receive updates about the page in their news feed. Page updates can include status updates, shared news articles, pictures, memes, or unique content. The name or profile pictures of people who like the page may be shown on the page or in ads about the page, giving it more legitimacy. According to the House Permanent Select Committee on Intelligence, the IRA created approximately 80,000 pieces of unique content to share within the 470 pages they created in this fashion (House Permanent Select Committee on Intelligence, 2018). To date, we do not specifically know what any of this unique content looked like, but we know that more than 126 million Americans were exposed to it. Therefore, our analysis of these Facebook advertisements helps us understand the IRA tactics to connect people to networks of disinformation content.

What specifically interests us is the emotion in the ad text, the role of affective wording in the ad text, and how emotion may be related to the ads' effectiveness in terms of engagement metrics. Was there a dominant affective tone in the Facebook ads? Was ad sentiment related to more engagement? Finally, were

---

<sup>4</sup> An impression with regard to Facebook ads is a measure of how often ads were on screen for an ad's target audience.



there changes over time in the ad sentiment? Did they appear to be linked in terms of sentiment to an election timeline? Although we know quite a bit about the ad targeting and the topics invoked in the ads, the emotional content or feelings evoked in the ad text seem to have been obscured by the discourse around fake news. The gaps in current studies of the ad content and presumed intentions led to our primary research questions:

*RQ1a: What is the valence of the emotional appeals in the Russian ads? How can we characterize them in terms of emotional content or sentiment?*

*RQ1b: Is there a relationship between sentiment and engagement, as measured by clicks and impressions?*

*RQ2: Is there a change over time in terms of sentiment in the IRA's Facebook ads? How might such change be linked to political campaigns in the United States?*

### Method



We analyzed the metadata in all of the 3,519 Facebook ads placed by the Russian IRA and released by Congress. By way of background, all 3,519 ads promoted one of the 470 IRA-created Facebook pages. Each advertisement was released as a two-page PDF. The first page included the metadata for the advertisement, which noted the ad identification; ad text; ad impressions (how often ads were on audience members' pages); ad clicks (number of time people actually clicked on an impression); ad landing page (often a fake group site); and ad targeting, which included targeted location, excluded connections, targeted age range, placement, and targeted interests (see Figure 1). The second page reproduced the actual ad. In Figure 1, our analysis focused on the sentiment in the ad text lines. After an initial cleaning process of deleting ads with inconclusive or erroneous metadata, we used 3,425 advertisements for analysis.

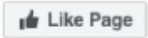
Ad ID 2774  
 Ad Text Heritage, not hate. The South will rise again! South United  
 Ad Landing Page <https://www.facebook.com/South-United-1777037362551238/>  
 Ad Targeting Location: United States  
 Excluded Connections: Exclude people who like South United  
 Age: 18 - 65+  
 Language: English (UK) or English (US)  
 Placements: News Feed on desktop computers, News Feed on mobile devices or Right column on desktop computers  
 People Who Match: Interests: Confederate Flag, Confederate States Army, American Civil War or Southern Pride  
 Ad Impressions 511,224  
 Ad Clicks 40,134  
 Ad Spend 78,357.33 RUB  
 Ad Creation Date 10/14/16 05:54:10 AM PDT

**Figure 1. Metadata example.**

To measure emotion in the text, we conducted a sentiment analysis of ad text in the data set. There are many different approaches to sentiment analysis, the most common focusing on document-level and/or individual word-level sentiment identification. Neither of these approaches was particularly effective for the current case because Facebook ad text generally occurs in one- to three-sentence units. The ads also generally incorporate an image. Therefore, we sought sentence-level analysis to study sentiment in Facebook ads.

Ad sentiment was evaluated using the *RSentiment* Package in the R development framework, a technique developed by combining word-level sentiment analysis with natural language processing techniques to effectively evaluate sentiment at or near the sentence level of granularity. *RSentiment* is an unassisted sentiment analysis technique that begins by assigning each word in a unit a sentiment score (positive, negative, very positive, very negative, or neutral). Word-level sentiment analysis is based on the conical Liu–Hu opinion lexicon (Liu, Hu, & Cheng, 2005). *RSentiment*'s natural language processing framework then evaluates the text unit for positive or negative sentiment or sarcasm and assigns a numerical score to indicate the unit's overall sentiment. The rating scale is continuous, with positive values indicating positive sentiment and negative values indicating negative sentiment. *RSentiment* assigns a score of 99 to any text unit that is likely to be sarcastic. If there are many positive and negative sentiments in the ad text, it is possible that *RSentiment* would code it as neutral; however, we did not observe that the ads typically contained mixed emotions. A comparative assessment of *RSentiment*'s accuracy found that it was more effective than other popular variants for social media and SMS-length units of analysis (Bose et al., 2017). Figures 2 and 3 illustrate the most negative and positive ads based on our sentiment analysis, and Figure 4 is an example of a sarcastic ad.


 **Being Patriotic**  
Sponsored · 

 Like Page

Boston police shot and killed a man wearing body armor and wielding an assault rifle who critically injured two officers responding to a domestic disturbance call late Wednesday, according to Police Commissioner William Evans.

A gun battle raged at an East Boston home as a suspect, Kirk Figueroa, 33, of East Boston, critically injured two Boston police officers late on October 12. He was then shot and killed by other officers who ran into the home to drag out their wounded col... [See More](#)

**ANOTHER GRUESOME ATTACK ON POLICE  
BY A BLM MOVEMENT ACTIVIST**



*Being Patriotic*

**OUR HEARTS ARE  
WITH THOSE 11 HEROES**

1.2K Reactions 82 Comments 377 Shares

**Figure 2. Most negative ad.**

 **Black Excellence**  
Sponsored · 



Leon Moss, 12, is Atlanta resident. His mother discovered Leon's talent when he was 9-year-old. She began to take him to local museums such as the Apex, the Sights and Sounds Black Cultural Museum and the Fernbank Museum of Natural History.  
To improve Leon's skills of painting he began to attend art classes of R. Gregory Christies. Non-profit organization the Southwest Atlanta Youth Business Organization (SWAYBO) also helped Leon in many ways.  
According to [blacknews.com](#), the ... [See More](#)



**LEON MOSS,  
12-YEAR-OLD ARTIST  
LAUNCHES INSPIRATIONAL  
"ART FOR KIDS" BUSINESS**

218 Reactions 14 Comments 61 Shares

 Like  Comment  Share

**Figure 3. Most positive ad.**



**Figure 4. Sarcasm example.**

For analytical purposes, we parsed the ad files and explored the sentiment in the ad text using RSentiment. Second, we sought to examine the relationship between ad sentiment and popularity—and hence the probability that people may have forwarded or shared the ad content and joined a Facebook group created by the IRA—by using the engagement measures of impressions and clicks, indicators present in the metadata for each ad. We note that 800 of the ads had no data regarding impressions or clicks, and we removed these from analyses investigating those two impact measures.

For the main analysis, we conducted a multiple linear regression analysis testing the effects of sentiment (a continuous variable representing degree of positive and negative), sarcasm (a binary sentiment variable), and duration of the advertisement on the number of clicks and impressions. The continuous sentiment scores calculated by RSentiment ranged from  $-21$  to  $21$  ( $M = 0.632$ ,  $SD = 2.997$ ). The raw scores assigned sarcasm a 99. These were recoded into a dichotomous categorical variable indicating whether the sentiment was sarcasm or not. The impressions and clicks variables had substantial outliers. To minimize this distributional problem, we calculated average impressions ( $M = 1165.9$ ,  $SD = 3787.62$ ) and clicks ( $M$


= 1067.9,  $SD = 299.26$ ) per sentiment score. The duration of the advertisement was calculated using the date the ads were created and ended (available in the metadata). The duration was calculated in the unit of seconds and then transformed into unit of days. There were 201 observations that had erroneous metadata (e.g., an earlier ad ending date than the ad starting date). For these and all others without the date information in the metadata, the duration was coded as zero.

Finally, to explore the third question, we developed a timeline to plot the ads' sentiment scores over time to investigate whether there were temporal trends of sentiment. Using the mean monthly sentiment scores, we plotted the swing of positive and negative emotions against their release timeline.

## Results

### *Ad Sentiment: Positive, Neutral, or Negative?*

To answer the first question, we plotted the frequencies of sentiment scores across the ads. The analysis in Figure 5 shows that outside of ads with no sentiment (value = 0), the majority (1,474) of the ads had positive sentiment (values  $\geq 1$ ). In fact, ads with positive sentiment were nearly double when compared with negative (789), neutral (725), or sarcastic (437) ads. The high number of neutral ads is also notable.



**Figure 5. Distribution of sentiment scores across the ads. Negative scores are negative sentiment, positive are positive sentiment.**

### Sentiment and Engagement

To investigate the relationship between the ads' sentiment and the number of impressions and clicks they received, we conducted two multiple regression analyses investigating effects of sentiment, advertisement duration, and sarcasm on impressions and clicks (calculated as average per sentiment scores). Table 1 summarizes the regression analysis results.

**Table 1. Multiple Regression Analysis Results of Sentiment, Sarcasm, and Ad Duration on the Advertisement Engagements.**

Variable	Multiple linear regression				Polynomial regression			
	<i>B</i>	<i>SE</i>	$\beta$	<i>t</i>	<i>B</i>	<i>SE</i>	$\beta$	<i>t</i>
Impressions per sentiment scores								
Sentiment	487.28	0.361	0.361***	24.259	939.53	0.016	0.654***	40.426
Duration	-8.93	0.015	-0.019	-1.269	-8.37	0.012	-0.018	-1.454
Sarcasm	-3563.57	0.045	-0.941***	-21.085	-3819.44	0.037	-1.08***	-27.411
Sentiment <sup>2</sup> (quadratic)					-48.96	0.003	-0.115***	-30.621
Sentiment <sup>3</sup> (cubic)					-3.95	0.001	-0.023***	-25.212
	$F(3, 3421) = 372.3^{***}, R^2 = .246$				$F(5, 3419) = 673.3^{***}, R^2 = .496$			
Clicks per sentiment scores								
Sentiment	30.27	0.015	0.284***	18.479	62.65	0.016	0.531***	33.300
Duration	-0.35	0.015	-0.009	-0.611	-0.37	0.012	-0.010	-0.816
Sarcasm	-287.36	0.046	-0.960***	-20.850	-321.69	0.036	-1.075***	-29.632
Sentiment <sup>2</sup> (quadratic)					-5.28	0.004	-0.150***	-40.341
Sentiment <sup>3</sup> (cubic)					-0.25	0.001	-0.019***	-20.575
	$F(3, 3421) = 279.6^{***}, R^2 = .197$				$F(5, 3419) = 711.8^{***}, R^2 = .51$			


\* $p < .05$ . \*\* $p < .01$ . \*\*\* $p < .001$ .

### Sentiment and Impressions

Our fitted multiple regression model was found to explain a substantial amount of variance in the impressions overall,  $F(3, 3421) = 372.3, R^2 = .246$ . Specifically, sentiment ( $\beta = 0.361, p < .001$ ) and sarcasm ( $\beta = -0.941, p < .001$ ) were found to be significant predictors of advertisement impressions. The duration of the advertisement was not a statistically significant predictor of advertisement impression ( $\beta = -0.019, p = .205$ ). In other words, as the sentiment increased by one unit, the expected average impression per sentiment scores increased by approximately 487.28 when all other variables were statistically held

constant. For sarcasm, the impressions were expected to be lower by  $-3563.57$  for advertisements with sarcastic sentiment compared with those that were not sarcastic when all other variables were controlled.

Although the conventional multiple linear regression predicted substantial variance in impressions, our visual inspection of the relationship between our focal predictor sentiment and impression indicated that the relationship could be nonlinear. Therefore, an additional analysis using polynomial terms was conducted. A model including the cubic polynomial term of the sentiment scores successfully explained 49.6% of the variance in the impressions,  $F(5, 3419) = 673.3$ ,  $R^2 = .496$ . Figure 6 displays the scatterplot and fitted line of the cubic polynomial model.




**Figure 6. Scatterplot and cubic polynomial regression of sentiment on impressions.**

### **Sentiment and Clicks**

The second multiple linear regression model predicting the number of clicks per sentiment score was found to explain a significant amount of variance overall,  $F(3, 3421) = 279.6$ ,  $R^2 = .197$ . Similar to the impressions, sentiment ( $\beta = 0.284$ ,  $p < .001$ ) and sarcasm ( $\beta = -0.960$ ,  $p < .001$ ) were found to be statistically significant predictors. The duration of the advertisements did not significantly predict the advertisement clicks per sentiment score ( $\beta = -0.009$ ,  $p = .541$ ). To elaborate, as the sentiment increased by one unit, the expected average clicks per sentiment score were expected to increase by approximately 30.27 when all other variables were statistically held constant. For sarcasm, the number of clicks was expected to be lower by  $-287.36$  for ads with sarcastic sentiment compared with those that were not sarcastic when all other variables were statistically controlled.




A preliminary visual examination hinted at similar concerns as with the previous analysis on impressions. Therefore, an identical cubic polynomial regression analysis was conducted. The model with the cubic polynomial terms of sentiment scores significantly explained 51% of the variance in the number of clicks,  $F(5, 3419) = 711.8$ ,  $R^2 = .51$ . Figure 7 displays the scatterplot and the fitted line of the cubic polynomial model.




**Figure 7. Scatterplot and cubic polynomial regression of sentiment on clicks.**

#### **Ad Sentiment Over Time**

Finally, to investigate our third research question about ad sentiment changes over time, we examined the distribution of sentiment at different time periods. Figure 8 plots sentiment scores over time and appears to show little dramatic change over time in ad sentiment. The blue line, representing a regression on sentiment across time, appears to vary little. However, a more granular examination (see Figure 9) of the monthly means of ad sentiment is more revealing. There, the negative swing of the ad's sentiment is visible in the period before the November 2016 election. September's ads were especially negative. The positive sentiment after the election is unmistakable. The gray area represents a confidence interval at 95%. We can also see that the positive ads garnered more clicks in general. The timeline does illustrate the variance over time, and it is tempting to conclude that some of it is deliberately in line with domestic political events.



**Figure 8. Mean sentiment over time.**



**Figure 9. Monthly mean sentiment over time.**

### Discussion and Conclusion

To briefly summarize the results, the majority of the ad texts were positive, followed by negative, neutral, and sarcastic ads. Ad text with positive sentiment garnered more impressions and clicks than did ads with neutral and negative sentiment. Finally, the ad text sentiment was more negative before November and there appears to be a swing upward immediately after the election.

Our first conclusion is that these ads' emotional content seems to be related to their being shared more broadly. The results underscore the power of "positive" messages (Berger & Milkman, 2012). We know from previous literature that social media content that evokes highly positive or negative physiological arousal is more viral, and that seems to be confirmed here. For Facebook specifically, posts with positive appeals, such as with humor and enthusiasm, are the most likely to be shared (Borah, 2016). One does not usually associate ads from politically motivated organizations with humor and uplifting messages, but in this particular case, both did occur, and the sentiment scores illustrate exactly that. Positive ads may have been used to grow the size of the Facebook groups the IRA was both creating and targeting, and subsequently these continued ads could have increased their legitimacy. (Some of the most popular Russian Facebook ads were for fake groups called Black Matters and Woke Blacks.)

We hypothesize that a tactic of the Russia IRA to use positive emotion in the ads was to help inculcate the in-group-out-group antagonism on which other researchers have commented. We know that

when social identities are made salient, people tend to put greater weight on their identity-relevant attitudes and favor their in-group. Using positive messages to prime people about their relevant in-group can enhance a sense of identity that then produces greater outrage when attacked. It may be that the IRA deployed positive sentiments for social groups, such as Blacks and Latinos, to make those identity categories more salient and later deployed negative sentiments by exposing them to attacks, for example. The net result could be polarization.

Our conclusions are limited in several important ways. First, any conclusion about the Russian IRA is incomplete without a full analysis of all of its social media content. Any attempt to outline a strategy must extend to the tactics used to raise membership for its Facebook pages as well as the IRA use of Twitter and other social media. Second, our sentiment analysis was restricted to the text that accompanied the Facebook but not the text in the ad picture. Although we would expect the text sentiment to be coherent, that may not always have been the case. Third, our inferences regarding strategic intentions behind the ads are necessarily somewhat speculative.

Despite these limitations, the findings in this study illustrate that the Russian disinformation campaign is more than made-up or "fake news." The integrated use of positive sentiment and negative sentiment may elicit the sort of emotional engagement that ultimately provokes concerns, perhaps about one's own social group and its representation and values in broader issues; then, over time and especially before the election, using negative sentiment could produce anger or disaffection, and even exacerbate existing social cleavages between groups. Our findings suggest therefore that combatting disinformation must go beyond flagging information sources as suspicious or false. Rather, the ads convey emotion and play to the strengths of the social media platform in encouraging quick bursts of engagement typical of "thinking fast." Emotional messages ultimately can contribute to feelings such as pride or anger. The Russian strategy, in many ways, may have worked to suggest certain ideas to different groups that they possibly already felt or believed were true (confirmation bias), playing off the emotional components of the ads and the platform that facilitate quick and unthinking reaction. Someone concerned about immigration may be more susceptible to advertisements that compare illegal immigrants with unlawful invaders. A red flag on the ad might not sway these individuals from feelings that their country is under threat. Indeed, Lewandowsky, Ecker, and Cook (2017) comment on "back-fire" effects of attempts that might accompany more rational debunking of ad content; emotional content is an armor against the rational, perhaps. The emotional content of the ads flies in the face of attempts to educate people to make rational decisions about such ads: Heuristic cognitive processing powerfully undercuts such an approach. To address this issue, social networking platforms themselves should have different standards for commercial and political issue advertising. Platforms might develop standards that evaluate sources as well as levels of emotion inherent in content.

Last, legislation might consider tightening the reach of political advertising of all flavors that has proliferated since Citizens United's legal decision which allowed for any group or individual, including foreign, to get involved in election campaigns with few campaign finance disclosure and reporting requirements (*Citizens United v. Federal Election Committee*, 2010). We now know that manipulated ads have been used widely by many players globally, including groups in the United States in state-level elections. Internationally, Sweden seems to have taken the lead in researching and recommending ways to counter

information-influence activities (Pamment, Nothhaft, Agardh-Twetman, & Fjallhed, 2018); however, more is known about the information-influence activities than what is known about countering them. The suggestions we offer are in and of themselves incomplete responses to address the threats of disinformation, but they do indicate that simply correcting "fake" or false information will do very little to protect against disinformation campaigns in the future.

### References

- Abelson, R. (1963). Computer simulation of hot cognition. In S. Tomkins & S. Messick (Eds.), *Computer simulation of personality* (pp. 277– 298). New York, NY: Wiley.
- Anderson, N., & Hubert, S. (1963). Effects of concomitant verbal recall on order effects in personality impression formation. *Journal of Memory and Language*, 2(5), 379–391. doi:10.1016/S0022-5371(63)80039-0
- Anstead, N., Magalhaes, J., Stupart, R., & Tambini, D. (2018, September). *Political advertising on Facebook: The case of the 2017 United Kingdom general election*. Paper presented at the European Consortium of Political Research Annual General Meeting, Hamburg, Germany.
- Bakir, V., & McStay, A. (2018). Fake news and the economy of emotions: Problems, causes, solutions. *Digital Journalism*, 6(2), 154–175. doi:10.1080/21670811.2017.1345645
- Bekafigo, M., & Pingley, A. (2015). Tweeting negative: Determinants of negative campaigning in the 2011 gubernatorial elections. *International Journal of E-Politics*, 6(1), 30–41. doi:10.4018/IJEP
- Benkler, Y., Faris, R., & Roberts, H. (2018). *Network propaganda: Manipulation, disinformation, and radicalization in American politics*. New York, NY: Oxford University Press.
- Berger, J., & Milkman, K. (2012). What makes online content viral? *Journal of Marketing Research*, 49(2), 192–205. doi:10.1509%2Fjmr.10.0353
- Borah, P. (2016). Political Facebook use: Campaign strategies used in 2008 and 2012 presidential elections. *Journal of Information Technology & Politics*, 13(4), 326–338.
- Borah, P., Fowler, E., & Ridout, T. (2018). Television vs. YouTube: Political advertising in the 2012 presidential election. *Journal of Information Technology & Politics*, 15(3), 230–244. doi:10.1080/19331681.2018.1476280
- Bose, S., Saha, U., Kar, D., Goswami, S., Nayak, A., & Chakrabarti, S. (2017). RSentiment: A tool to extract meaningful insights from textual reviews. In S. Satapathy, V. Bhateja, S. Udgata, & P. Pattnaik (Eds.), *Proceedings of the Fifth International Conference on Frontiers in Intelligent*

- Computing: Theory and applications. Advances in intelligent systems and computing* (Vol. 516, 259–268). Singapore: Springer.
- Brader, T. (2006). *Campaigning for hearts and minds: How emotional appeals in political ads work*. Chicago, IL: University of Chicago Press.
- Brady, W., Wills, J., Jost, J., Tucker, J., & Van Bavel, J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, *114*(28), 7313–7318. doi:10.1073/pnas.1618923114
- Buck, R., Chaudhuri, A., Georgson, M., & Kowta, S. (1995). Conceptualizing and operationalizing affect, reason, and involvement in persuasion: The ARI model and the CASC scale. *Advances in Consumer Research*, *22*, 440–447. Retrieved from <https://www.acrwebsite.org/volumes/7785/volumes/v22/NA-22>
- Carver, C., & Harmon-Jones, E. (2009). Anger is an approach-related affect: Evidence and implications. *Psychological Bulletin*, *135*(2), 183–204.
- Chaiken, S. (1980). Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology*, *39*(5), 752–766. doi:10.1037/0022-3514.39.5.752
- Chaiken, S., & Ledgerwood, A. (2012). A theory of heuristic and systematic information processing. In P. Van Lange, A. Kruglanski, & T. Higgins (Eds.), *Handbook of theories of social psychology* (Vol. 1, pp. 246–266). Thousand Oaks, CA: SAGE Publications. doi:10.4135/9781446249215
- Chaudhuri, A., & Buck, R. (1995). Media differences in rational and emotional responses to advertising. *Journal of Broadcasting & Electronic Media*, *39*(1), 109–125. doi:10.1080/08838159509364291
- Citizens United v. Federal Election Commission, 558 U.S. 310 (2010).
- Clore, G., & Isbell, L. (2001). Emotions as virtue and vice. In J. Kuklinski & D. Chong (Eds.), *Citizens and politics: Perspectives from political psychology* (pp. 103–126). New York, NY: Cambridge University Press.
- Cyr, D., Head, M., Lim, E., & Stibe, A. (2018). Using the elaboration likelihood model to examine online persuasion through website design. *Information & Management*, *55*(7), 807–821. doi:10.1016/j.im.2018.03.009
- Dahlgren, P. (2018). Public sphere participation online: The ambiguities of affect. *International Journal of Communication*, *12*, 2052–2070. Retrieved from <https://ijoc.org/index.php/ijoc/article/view/6786>

Druckman, J., Kifer, M., & Parkin, M. (2010). Timeless strategy meets new medium: Going negative on congressional campaign Web sites, 2002–2006. *Political Communication*, 27(1), 88–103. doi:10.1080/10584600903502607

Dwoskin, E., & Romm, T. (2018, July 31). Facebook says it has uncovered a coordinated disinformation operation ahead of the 2018 midterm elections. *The Washington Post*. Retrieved from [https://www.washingtonpost.com/technology/2018/07/31/facebook-says-it-has-uncovered-coordinated-disinformation-operation-ahead-midterm-elections/?utm\\_term=.5190c407999b](https://www.washingtonpost.com/technology/2018/07/31/facebook-says-it-has-uncovered-coordinated-disinformation-operation-ahead-midterm-elections/?utm_term=.5190c407999b)

Eagly, A., & Chaiken, S. (1993). *The psychology of attitudes*. San Diego, CA: Harcourt Brace Jovanovich.

Elliott, A. (2008). *Concepts of the self*. Cambridge, UK: Polity Press.

Facebook. (n.d.). *Facebook ad formats for different goals*. Retrieved from [https://www.facebook.com/business/ads/adformats?ref=ads\\_guide/](https://www.facebook.com/business/ads/adformats?ref=ads_guide/)

Fowler, E., Ridout, T., & Franz, M. (2016). Political advertising in 2016: The presidential election as outlier? *The Forum*, 14(4), 445–469. doi:10.1515/for-2016-0040

Gelfert, A. (2018). Fake news: A definition. *Informal Logic*, 38(1), 84–117. doi:10.22329/il.v38i1.5068

Gervais, B. (2017). More than mimicry? The role of anger in uncivil reactions to elite political incivility. *International Journal of Public Opinion Research*, 29, 384–405. doi:10.1093/ijpor/edw010

Goldman, A., Barnes, J., Haberman, M., & Fandos, N. (2020, February 20). Lawmakers are warned that Russia is meddling to re-elect Trump. *The New York Times*. Retrieved from <https://www.nytimes.com/2020/02/20/us/politics/russian-interference-trump-democrats.html>

Guynn, J., Weise, E., & Kelly, E. (2018, May 11). Release of Facebook ads shows pattern. *USA Today*. Retrieved from <https://www.pressreader.com/usa/usa-today-us-edition/20180511/281487866989070>

Groshek, J., & Koc-Michalska, K. (2017). Helping populism win? Social media use, filter bubbles, and support for populist presidential candidates in the 2016 U.S. election campaign. *Information, Communication & Society*, 20(9), 1389–1407. doi:10.1080/1369118X.2017.1329334

Hogg, M., Terry, D., & White, K. (1995). A tale of two theories: A critical comparison of identity theory with social identity theory. *Social Psychology Quarterly*, 58(4), 255–269. Retrieved from [www.jstor.org/stable/2787127](http://www.jstor.org/stable/2787127)

House Permanent Select Committee on Intelligence. (2018). *Exposing Russia's efforts to sow discord online: The internet research agency and advertisements* [White Paper]. Retrieved from <https://intelligence.house.gov/social-media-content/>

- Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., & Westwood, S. (2019). The origins and consequences of affective polarization in the United States. *Annual Review of Political Science*, 22, 129–146. doi:10.1146/annurev-polisci-051117-073034
- Iyengar, S., Sood, G., & Lelkes, Y. (2012). Affect, not ideology: A social identity perspective on polarization. *Public Opinion Quarterly*, 76(3), 405–431. doi:10.1093/poq/nfs038
- Jackson, D. (2017, October 17). *Distinguishing disinformation from propaganda, misinformation, and "fake news."* Washington, DC: National Endowment for Democracy. Retrieved from <https://www.ned.org/issue-brief-distinguishing-disinformation-from-propaganda-misinformation-and-fake-news/>
- Jamieson, K. (2018). *Cyberwar: How Russian hackers and trolls helped elect a president*. New York, NY: Oxford University Press.
- Kahneman, D. (2011). *Thinking fast and slow*. London, UK: Penguin Books.
- Katz, E., & Lazarsfeld, P. (1955). *Personal influence: The part played by people in the flow of mass communications*. London, UK: Transaction Publishers.
- Kelly, M. (2019, April 29). New data will give insight into Facebook's influence on elections. *The Verge*. Retrieved from <https://www.theverge.com/2019/4/29/18522473/facebook-election-influence-meddling-research-data-democracy>
- Kim, Y. (2018). Uncover: Strategies and tactics of Russian interference in U.S. elections. Report from Project DATA. Retrieved from [https://journalism.wisc.edu/wp-content/blogs.dir/41/files/2018/09/Uncover.Kim\\_.v.5.0905181.pdf](https://journalism.wisc.edu/wp-content/blogs.dir/41/files/2018/09/Uncover.Kim_.v.5.0905181.pdf)
- Kim, Y., Hsu, J., Neiman, D., Kou, C., Bankston, L., Kim, S., & Raskutti, G. (2018). The stealth media? Groups and targets behind divisive issue campaigns on Facebook. *Political Communication*, 35(4), 1–29. doi:10.1080/10584609.2018.1476425
- Kramer, A., Guillory, J., & Hancock, J. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academic of Sciences USA*, 111(29), 8788–8790. doi:10.1073/pnas.1320040111
- Kreiss, D., & Howard, P. (2010). New challenges to political privacy: Lessons from the first U.S. presidential race in the Web 2.0 era. *International Journal of Communication*, 4, 1032–1050. <https://ijoc.org/index.php/ijoc>
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480–498. doi:10.1037/0033-2909.108.3.480



- Lasswell, H. D. (1927). *Propaganda technique in the World War*. New York, NY: Ravenio Books.
- Lewandowsky, S., Ecker, U., & Cook, J. (2017). Beyond misinformation: Understanding and coping with the "post-truth" era. *Journal of Applied Research in Memory and Cognition*, 6(4), 353–369. doi:10.1016/j.jarmac.2017.07.008
- Liu, B., Hu, M., & Cheng, J. (2005). Opinion observer: Analyzing and comparing opinions on the Web. In *Proceedings of the 14th International Conference on the World Wide Web* (pp. 342–351). New York, NY: Association for Computing Machinery. doi:10.1145/1060745.1060797
- Lodge, M., & Taber, C. (2000). Three steps toward a theory of motivated political reasoning. In A. Lupia, M. McCubbins, & S. Popkin (Eds.), *Elements of reason, cognition, choice, and the bounds of rationality* (pp. 183–213). London, UK: Cambridge University Press.
- Marcus, G. (2002). *Sentimental citizen: Emotion in democratic politics*. University Park: Pennsylvania State University Press.
- Matsa, E., & Shearer, E. (2018). *News use across social media platforms 2018*. Washington, DC: Pew Research Center. Retrieved from <http://www.journalism.org/2018/09/10/news-use-across-social-media-platforms-2018>
- Mueller, R. (2019, March). *Report on the investigation into Russian interference in the 2016 presidential election*. Washington, DC: U.S. Department of Justice. Retrieved from <https://www.documentcloud.org/documents/5955118-The-Mueller-Report.html>
- Nyhan, B., & Reifler, J. (2010). When corrections fail: The persistence of political misperceptions. *Political Behavior*, 32(2), 303–330. doi:10.1007/s11109-010-9112-2
- Pamment, J., Nothhaft, H., Agardh-Twetman, H., & Fjallhed, A. (2018). *Countering information influence activities: The state of the art*. Karlstad, Sweden: Swedish Civil Contingencies Agency. Retrieved from <https://www.msb.se/RibData/Filer/pdf/28697.pdf>
- Penzenstadler, N., Heath, B., & Guynn, J. (2018, May 11). We read every one of the 3,517 Facebook ads bought by Russians. Here's what we found. *USA Today*. Retrieved from <https://www.usatoday.com/story/news/2018/05/11/what-we-found-facebook-ads-russians-accused-election-meddling/602319002/>
- Petty, R. E., & Cacioppo, J. T. (1984). The effects of involvement on responses to argument quantity and quality: Central and peripheral routes to persuasion. *Journal of Personality and Social Psychology*, 46(1), 69–81. <https://doi.org/10.1037/0022-3514.46.1.69>

- Pew Research Center. (2016). *The partisan divide on political values grows even wider*. Washington, DC: Author. Retrieved from <http://www.people-press.org/2017/10/05/the-partisan-divide-on-political-values-grows-even-wider/>
- Pi, S., Chou, C., & Liao, H. (2013). A study of Facebook groups members' knowledge sharing. *Computers in Human Behavior, 29*(5), 1971–1979. doi:10.1016/j.chb.2013.04.019
- Placencia, M., & Lower, A. (2013). Your kids are so stinkin' cute! :-): Complimenting behavior on Facebook among family and friends. *Intercultural Pragmatics, 10*(4), 617–646. doi:10.1515/ip-2013-0029
- Popkin, S. (1994). *The reasoning voter: Communication and persuasion in presidential campaigns*. Chicago, IL: University of Chicago Press.
- Ribeiro, F., Ufmg, U., Henrique, L., Gummadi, K., & Redmiles, E. (2018). On microtargeting socially divisive ads: A case study of Russia-linked ad campaigns on Facebook. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 140–149). New York, NY: Association for Computing Machinery. doi:10.1145/3287560.3287580
- Roose, K. (2018, December 17). Social media's forever war. *The New York Times*. Retrieved from <https://www.nytimes.com/2018/12/17/technology/social-media-russia-interference.html>
- Schwarz, N. (2011). Feelings-as-information theory. In P. Van Lange, A. Kruglanski, & T. Higgins (Eds.), *Handbook of theories of social psychology* (Vol. 1, pp. 289–308). Thousand Oaks, CA: SAGE Publications. doi:10.4135/9781446249215
- Silverman, C., & Alexander, L. (2016, November 3). How teens in the Balkans are duping Trump supporters with fake news. *BuzzFeed News*. Retrieved from <https://www.buzzfeednews.com/article/craigsilverman/how-macedonia-became-a-global-hub-for-pro-trump-misinfo>
- Stanley, J. (2015). *How propaganda works*. Princeton, NJ: Princeton University Press.
- Taber, C., & Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science, 50*(3), 755–769. doi:10.1111/j.1540-5907.2006.00214.x
- Tajfel, H. (1969). Cognitive aspects of prejudice. *Journal of Biosocial Science, 1*(Suppl. 1), 173–191. doi:10.1017/S0021932000023336
- Tajfel, H., & Turner, J. (1979). An integrative theory of intergroup conflict. In M. Hatch & M. Schultz (Eds.), *Organizational identity: A reader* (pp. 56–65). New York, NY: Oxford University Press.

- Tannenbaum, M., Hepler, J., Zimmerman, R., Saul, L., Jacobs, S., Wilson, K., & Albarracín, D. (2015). Appealing to fear: A meta-analysis of fear appeal effectiveness and theories. *Psychological Bulletin, 141*(6), 1178–1204. doi:10.1037/a0039729
- Thorson, K., Vraga, E., & Ekdale, B. (2010). Credibility in context: How uncivil online commentary affects news credibility. *Mass Communication and Society, 13*(3), 289–313. doi:10.1080/15205430903225571
- Van Duyn, E., & Collier, J. (2018). Priming and fake news: The effects of elite discourse on evaluations of news media. *Mass Communication and Society, 22*, 1–20. doi:10.1080/15205436.2018.1511807
- Weeks, B. (2015). Emotions, partisanship, and misperceptions: How anger and anxiety moderate the effect of partisan bias on susceptibility to political misinformation: Emotions and misperceptions. *Journal of Communication, 65*(4), 699–719. doi:10.1111/jcom.12164